

LIBRARY
OF THE
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY



28
414
393-
69

UTILITY THEORY APPLIED TO MEDICAL
DIAGNOSIS AND TREATMENT

by

NORMAN E. BETAQUE, JR.

393-69

May 1969



UTILITY THEORY APPLIED TO MEDICAL
DIAGNOSIS AND TREATMENT

by

NORMAN E. BETAQUE, JR.

393-69

May 1969

UTILITY THEORY APPLIED TO MEDICAL
DIAGNOSIS AND TREATMENT

by

NORMAN ELLSWORTH BETAQUE JR.

Submitted to the Department of Electrical Engineering on 20 May 1969 in partial fulfillment of the requirements for the Degrees of Master of Science and Electrical Engineer.

ABSTRACT

Attempts by others to model medical diagnosis and treatment as decision making under risk have found utility assessment to be a major stumbling block. Some have avoided the problem by using bivalued utilities. Others have conjectured about the difficulty of assessing utilities over medical costs and consequences, but have not seriously tried to make the assessments.

This thesis reports an attempt to implement a utility theory model of the diagnosis and treatment of a serious medical condition, acute renal failure. Doctors provided subjective estimates of the required conditional probabilities and assessed utilities over their preferences for tests, treatments and consequences. The model was then tested on twenty-eight hypothetical medical cases. The goal was to duplicate the decisions (but not necessarily the procedures leading to these decisions) made by qualified doctors on the cases.

The model's success in duplicating doctors' decisions on about ninety per cent of the hypothetical cases indicates that utility theory is not only a convenient structure for theoretically describing diagnosis and treatment, but that it is potentially a practical tool for analyzing such decision problems.

THESIS SUPERVISOR: G. Anthony Gorry

TITLE: Assistant Professor of Management

ACKNOWLEDGMENTS

I would like to express my appreciation to Professor G. Anthony Gorry for the guidance he has given me throughout this project. His interest and enthusiasm were a great encouragement. I also owe my gratitude to Dr. Jerome Kassirer and Dr. William Schwartz of the Pratt Diagnostic Clinic, New England Medical Center, for their generous contributions of time, information, and advice.

Finally, my special thanks to my wife Arleen and daughter Karen. Their confidence and co-operation were invaluable.

TABLE OF CONTENTS

	<u>Page</u>
ABSTRACT	ii
ACKNOWLEDGMENTS	iii
CHAPTER 1 - INTRODUCTION	1
1.1 The Objective	1
1.2 The Decision Tree	4
1.3 Utility Theory	9
CHAPTER 2 - LITERATURE SURVEY	14
CHAPTER 3 - DEVELOPMENT OF THE MODEL	17
3.1 General Approach	17
3.2 The Model	22
3.3 The Probabilities	30
3.4 The Utilities	34
CHAPTER 4 - IMPLEMENTATION OF THE MODEL	37
4.1 The Methodology	37
4.2 The First Attempt at Utility Assessment . .	40
4.3 The Second Attempt at Utility Assessment .	49
4.4 Utility Assessment by Dr. # 2	57
CHAPTER 5 - DISCUSSION	64
BIBLIOGRAPHY	71
APPENDIX 1 - CONDITIONAL PROBABILITIES OF CONSEQUENCES . . .	72
APPENDIX 2 - CONDITIONAL PROBABILITIES OF TEST RESULTS . . .	76

LIST OF TABLES

	<u>Page</u>
TABLE 1 - DISEASES	24
TABLE 2 - TREATMENTS AND CONSEQUENCES	25
TABLE 3 - TESTS	26
TABLE 4 - TEST RESULTS	27
TABLE 5 - HYPOTHETICAL CASES	38
TABLE 6 - INDIFFERENCE LOTTERIES (Dr. # 1 - First Attempt)	42
TABLE 7 - INDIVIDUAL UTILITY SCALES (Dr. # 1 - First Attempt)	43
TABLE 8 - UTILITY SCALE (Dr. # 1 - First Attempt) . . .	46
TABLE 9 - UTILITIES (Second Attempt)	50
TABLE 10 - DECISIONS ON HYPOTHETICAL CASES (Dr. # 1 - Initial Utilities)	52
TABLE 11 - DECISIONS ON THE HYPOTHETICAL CASES (Dr. # 1 - Adjusted Utilities)	55
TABLE 12 - UTILITY SCALE (Dr. # 2)	61
TABLE 13 - DECISIONS ON THE HYPOTHETICAL CASES (Dr. # 2)	63

CHAPTER I

INTRODUCTION

1.1 The Objective

The ability of computers to store large quantities of data, perform logical operations on the data, and provide rapid access to stored information has prompted research in their use for improving medical diagnosis and treatment. Part of that research has been directed towards the development of new techniques which take advantage of computers to provide better monitoring of the patient or better analysis of some test or treatment. Another direction of the research has been towards helping the doctor make better use of existing techniques. This project is a step in the latter direction.

Medical knowledge is expanding rapidly, and it is difficult for the average doctor to keep himself abreast of new developments. As a result, there is not only a trend toward specialization, but there is a growing gap between a specialist's knowledge and his ability to pass that knowledge on to other doctors who are less experienced in his field. Traditional means of communicating medical knowledge, such as seminars, publications, and conferences, can illustrate techniques or state statistical results of experiments. But, it is difficult for doctors to communicate judgments or experience through these kinds of media. It is current practice, when such judgment or experience is required, to either send the patient to a specialist or enlist a specialist as a

consultant on the case. Of course, there are not enough specialists in any field to permit one of them to be personally involved in every case needing their specialty. Nor do all doctors realize when a case has progressed beyond their competence and a specialist is required. As a result, many patients do not have the services of a specialist when his experience and judgment are badly needed.

The motivation for developing a computer-aided diagnosis-treatment system is to help fill the gap between the specialist and the average doctor, and to make some of the experience and judgment of the specialist readily available. In effect, a computer could, at least in part, play the same role as a consultant. It may not be quite as good as the specialist, himself, but it could greatly aid many doctors. Furthermore, computer facilities can be made much more accessible for general use than can a specialist.

The basic problem in trying to develop a computer-aided system for diagnosis and treatment is that doctors, in performing such tasks, rely on education, experience, preferences, intuition or judgment — all qualities of the individual doctor. A computer, however, can manipulate only quantitative data and requires a mathematical description of the process. One approach to describing diagnosis and treatment mathematically is to model it as a sequence of decisions made under conditions of risk. That is, at the time a decision is made, it is not known for sure which one of a possible set of consequences will occur. This type of analysis of a decision situation is commonly referred to as statis-

tical decision theory, decision analysis, or, as in this paper, utility theory.

This paper reports efforts to use utility theory to model the diagnosis and treatment of a specific medical syndrome, acute renal failure. The objective was to duplicate the decisions made by a doctor (but not necessarily the process employed by the doctor to reach those decisions) on a number of hypothetical cases. The primary interest of the project was to experience the practical problems of applying utility theory to medical diagnostic-treatment problems and to determine if the theory is a reasonable approach to such problems.

1.2 The Decision Tree

The structuring of a decision problem for analysis initially involves the identification of available courses of action, their possible consequences, and the conditions under which those consequences might occur. Consider, for example, the diagnosis-treatment problem diagrammed by the decision tree in Figure 1. There are two available courses of action: treat for disease A or treat for disease B. There are three possible consequences: the patient either is cured, dies, or his condition remains unchanged. The conditions under which each of the consequences might occur is dependent, not only on the course of action chosen, but on some unknown state of nature, in this case the particular disease the patient actually has. In this simple example, then, it is the unknown state of nature that introduces the element of risk into the problem. In real medical problems, of course, the consequences themselves may not be deterministic; the best treatment for a known disease may not guarantee a specific consequence. So, there is an element of risk even when the state of nature is known. A decision tree can be easily modified to include these situations by adding branches representing the additional consequences to appropriate nodes. Further extension of the problem can include provisions for improving the knowledge about the unknown state of nature by performing tests or experiments. In medicine, such tests might range from physical examination to major surgery. Whatever the case may be, the primary purpose of a test is to gain knowledge about the state of nature.

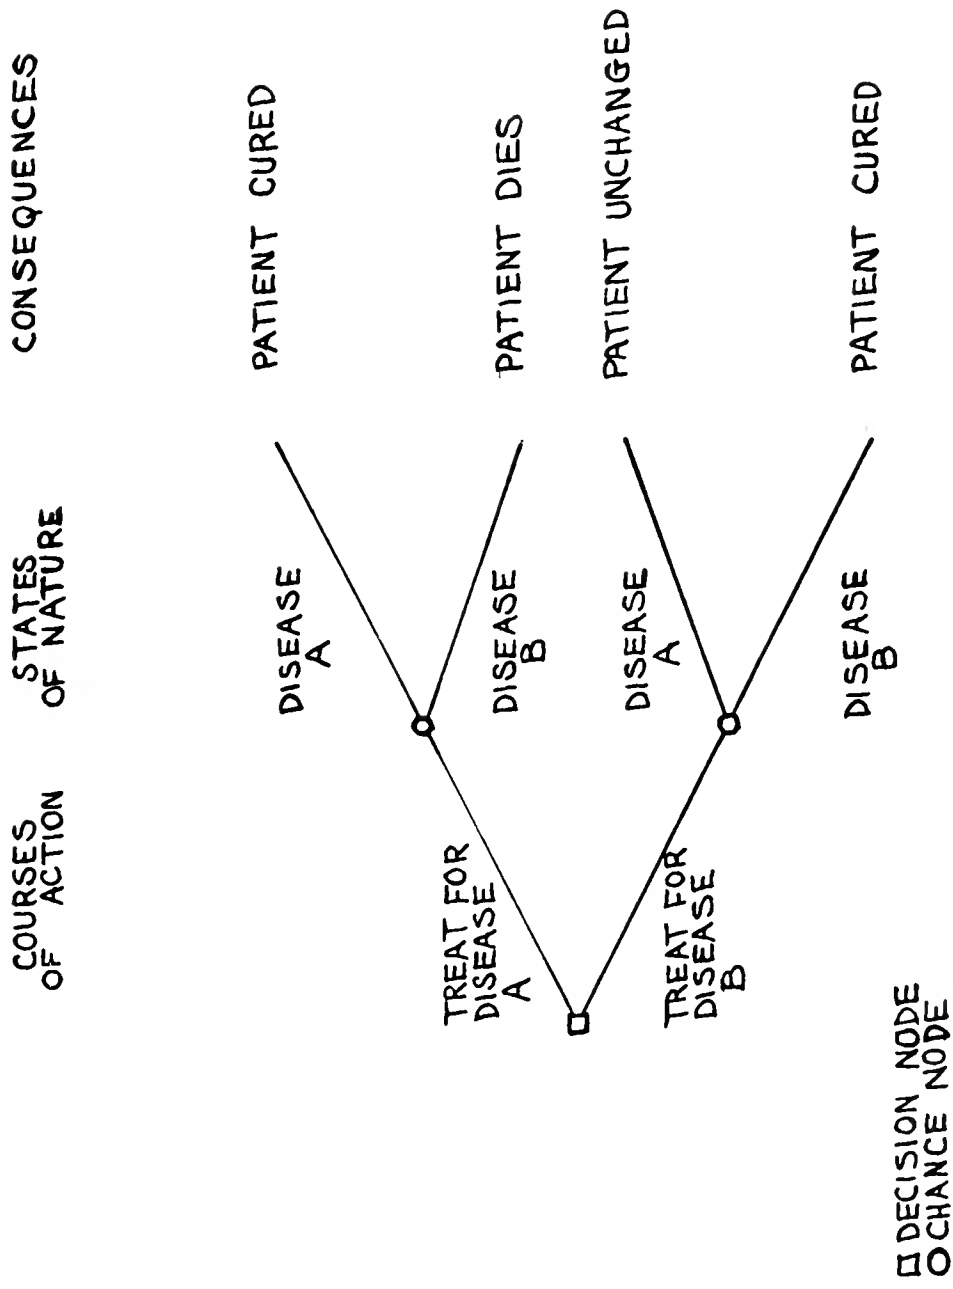


FIGURE I

Figure 2 illustrates the expanded decision tree.

Several aspects of the expanded problem are noteworthy. First, whether or not to perform a test is itself a decision having uncertain results. Obviously, if the results of a test are precisely predictable, it is senseless to perform the test. (If a test offers only therapeutic benefits it can be considered a treatment.) Note also that the ultimate consequences now include the costs of the tests. In some situations, the test costs are so small in relation to other consequences that they can safely be ignored; in others, however, the risk, patient discomfort, time, or inconvenience associated with a potentially useful test may lead a doctor to either carry out treatment without testing, or to give some other less informative, but less costly test. It should also be realized that the performance of a test will change the likelihood of a specific consequence occurring, as this is the whole purpose of testing: to reduce the risk of treatment by improving the knowledge about the unknown disease. It would seem desirable to conduct as many tests as possible to learn the true state of nature. Balanced against this desire, however, is the aggregate cost of testing, where cost may refer to time, pain, risk, and other factors as well as monetary expenses. At some point, the value of the knowledge attainable by further testing is exceeded by the costs. Hence it may be best to make the choice of treatments even though the actual disease is still, to some degree, uncertain.

It is in this process of weighing the merits of further testing,




FIGURE 2

re-evaluating the present state of knowledge, and finally choosing a course of action that a doctor invokes his own preferences and judgments. And it is these judgments and preferences which must be quantified to facilitate application of a computer to the decision problem. Utility theory offers a means of quantifying such qualitative information and provides an unambiguous criteria for selecting courses of action consistent with these preferences and judgments.

1.3 Utility Theory

Utility theory is an axiomatic theory of decision making under risk. It is a prescriptive theory; it specifies the course of action a decision maker should choose to be consistent with his preferences and judgments. In short, it states that if a decision maker accepts the axioms of the theory, he should choose that course of action which maximizes his expected utility, utility being a numerical scale over the decision maker's preferences.

The axioms of utility theory are fundamental assumptions which form the basis of the theory and from which can be derived all other aspects of the theory. On the surface, it appears that any man trying to act rationally would not only accept the axioms but would like to make all his decisions consistent with them. In reality, they are not as universally attractive as would first appear. In fact, there are several sets of axioms leading to similar variations of utility theory. Luce and Raiffa⁽⁶⁾ discuss some of the psychological implications of accepting a utility theory as well as the detailed mathematical statement of the axioms. Presented here is one set of axioms, or assumptions, taken from Chernoff and Moses⁽¹⁾.

Assumption 1. (Ordering of Alternatives) An individual faced with two prospects, A and B, will be able to decide whether he prefers A to B, B to A, or is indifferent between them.

Assumption 2. (Transitivity) If prospect A is preferred to B, and B is preferred to C, then A is preferred to C.

Assumption 3. If prospect A is preferred to prospect B, which is in turn preferred to C, then there is a mixture* of A and C which is preferred to B, and a mixture of A and C to which B is preferred.

Assumption 4. If prospect A is preferred to prospect B, and C is another prospect, then any mixture of A and C will be preferred to the same mixture of B and C.

These four assumptions are sufficient to establish a linear scale, called a utility scale, over the prospects. That is, to each prospect corresponds a number called the utility of that prospect. This utility scale has the following properties:

1. The utility of prospect A is greater than that of prospect B if and only if A is preferred to B.

2. If B is a prospect where, with probability p, the individual faces prospect A and with probability 1-p he faces prospect C, then

$$\text{Utility of B} = p(\text{Utility of A}) + (1-p)(\text{Utility of C})$$

Property 1 merely says that more preferred prospects have higher utilities. Property 2 states that if prospect B can be considered a lottery offering a p chance at prospect A and a 1-p chance at prospect C, then the utility of B is the expected utility of the lottery.

The computations made in applying utility theory to a decision problem stem directly from these two properties. First, it is necessary to assess utilities over all the consequences being considered.

* The nature of this mixture will be discussed later.

(Exactly how to do this is discussed in section 3.4.) Next, if the likelihood of prospects occurring can be expressed as probabilities, it is possible, by working backwards through the decision tree and invoking Property 2, to determine the expected utility of each course of action at each decision node. According to Property 1, then, the decision maker should choose, at each decision node, the course of action promising the highest expected utility.

To illustrate this procedure, suppose that a doctor faced with the simple decision problem described by Figure 1 has assessed the following utilities over the consequences.

Utility of patient dying	= 0
Utility of patient's condition not changing	= 10
Utility of patient's being cured	= 20

Suppose, further, that he has been able to express his judgment of the likelihood of prospects occurring in terms of the following probabilities.

Probability that patient has disease A	= .6
Probability that patient has disease B	= .4

The utilities of each of the available courses of action are therefore:

Utility of treating for disease A is	$.6(20) + .4(0) = 12$
Utility of treating for disease B is	$.6(10) + .4(20) = 14$

Therefore, the decision should be to treat for disease B, since that course of action promises the highest expected utility.

As the decision problem becomes more complex, the basic procedures

for analyzing the problem remain the same. The computation starts at the end of the decision tree and progresses backwards, step by step, through the sequence of decisions, determining the best course of action at each decision node. The task of judging the likelihood of prospects occurring, however, becomes increasingly difficult, for, now, the initial, or a priori, probabilities of the diseases are insufficient to express the judgments involved. To adequately express these judgments, the probabilities must be conditioned on all decisions and test results attained up to that juncture in the decision tree. For example, in Figure 2, to determine the expected utility of treating for disease A after test result A has been received, it is necessary to know the conditional probability of each of the diseases given that the test has been done and has yielded result A. If more than one test is done, the probabilities must be conditioned on each of the possible test results. Essentially, this is the process of updating the knowledge of the unknown state of nature as a result of information gained through testing.

Although this is qualitatively the type of inference a decision maker considers in evaluating the significance of tests, the direct expression of such judgments as probabilities is usually an extremely difficult, if not impossible, task. Fortunately, these required probabilities can be readily computed from conceptually easier conditional probabilities using Bayes Theorem. Because Warner et al.⁽¹⁰⁾ and Gorry and Barnett⁽⁴⁾ have discussed Bayes Theorem and its application as an

inference function in medical diagnosis, it will suffice to state here that it provides a convenient method for determining the probability that a patient has a particular disease given the results of a series of tests. The only judgments necessary are the conditional probabilities of the test results given each disease and the a priori probabilities of each disease. Though perhaps numerous, these probabilities can usually be estimated from past experience or historical data.

Thus, utility theory, with the computational aid of Bayes Theorem, provides a formal structure for quantitatively analyzing a decision problem and an unambiguous criteria for choosing among courses of action. The only quantitative information necessary for its application is the decision maker's preferences expressed as utilities and his judgments expressed as simple conditional probabilities.

CHAPTER 2

LITERATURE SURVEY

Several studies have been directed towards evaluating the potential of Bayes Theorem as an inference function for medical diagnosis. Warner et al.⁽¹⁰⁾ applied a model based on Bayes Theorem to the diagnosis of congenital heart disease. Overall and Williams⁽⁷⁾ did similarly with conditions of thyroid function, as did Lodwick et al.⁽³⁾ with primary bone tumors. The approach in these studies was to accumulate statistical data on the incidence of signs and symptoms in the diseases being considered. The particular signs and symptoms observed in any one patient were then transformed using the data and Bayes Theorem into the probability of the patient having each of the diseases. As all observations of the patient were made before the model was invoked, the model served solely to infer a diagnosis from a given set of signs and symptoms. Then studies all reported diagnostic results favorably comparable to those achieved by well-qualified clinicians.

Gorry⁽³⁾ further developed the model of medical diagnosis along the concepts of utility theory by using the model to select the signs and symptoms to be observed as well as to infer a diagnosis. His model was similar to that described earlier in Chapter 1; however, his only courses of action were tests and diagnoses, and his only consequences were a correct diagnosis or a misdiagnosis. In addition, rather than attempt utility assessment over the test costs or misdiagnoses, he

arbitrarily assigned equal utilities to all test costs and equal utilities to all misdiagnoses. Thus, although his decision policy was to minimize expected utility, he was, in effect, minimizing the number of tests required to achieve the most probable diagnosis. Using Warner's data on congenital heart disease, Gorry demonstrated that by using Bayes Theorem to select tests as well as evaluate test results, the number of tests required to achieve a good diagnosis could be drastically reduced. He also demonstrated that by varying the utilities of test costs or consequences he could change the optimum sequence of testing.

Rubel⁽⁹⁾ used utility theory to model three problems involving medical diagnosis and treatment: a patient complaining of a sore throat, one suffering from a possible kidney malfunction, and a medical officer of the Food and Drug Administration trying to decide whether or not to license a new drug. Rubel concentrated on developing the theoretical aspects of the model and hypothesizing the possible difficulties in implementing them. He did not gather data or elicit probabilities or utilities from doctors. He did, however, discuss several aspects of assessing utilities over consequences. His suggested approach was to identify various attributes of the consequences, have the patient assess utilities over each of the attributes individually, then try to obtain some relationship or trade-off among the attributes to enable the preference for any consequence to be expressed as a single numeraire. Also proposed was a technique of graphically presenting the best courses

of action as functions of the separate attributes; the doctor was then to use the graph and his evaluation of the patient's trade offs to determine the best policy.

CHAPTER 3

DEVELOPMENT OF THE MODEL

3.1 General Approach

As indicated in Chapter 2, the theoretical implications of applying utility theory to the analysis of medical diagnosis-treatment problems have been considered by several authors. Likewise, the use of Bayes Theorem as an inference function for selecting and processing information pertinent to medical diagnosis has also been previously examined. The practical aspects of assessing utilities, however, and applying the full potential of utility theory as a basis for a computer-aided diagnostic-treatment system have not been explored. It is to this task that the present project was directed.

The approach was to model the diagnosis and treatment of a specific syndrome (acute renal failure) using the analysis techniques appropriate to the application of utility theory. Effort was made to keep the model as uncomplicated as possible. The goal was merely to determine if utility theory offered a reasonable approach to analyzing the problem, rather than to develop an accurate, working, diagnostic-treatment system. In fact, only those tests thought to have non-trivial costs or which were useful in very special situations were considered. Routine tests and questions pertaining to medical history or physical examination were not included in the project. Further simplification of the model was achieved by consolidating similar treatments or conse-

quences into one category. For instance, a drug can be administered in various doses; in this project, only one dose of each drug, usually large, was considered. Likewise, rather than specify in detail the possible results of treatment, all consequences were placed into one of several general categories. Thus the emphasis was on determining if, within the context of a simplified problem, utility theory provided a practical basis of describing medical diagnosis and treatment.

A major aspect of this project was the use of a computer* to perform all the calculations required in using the model. After the utilities and probabilities had been obtained from a doctor, the computer was used to rapidly determine the expected utility of each course of action and to specify the most desirable course of action. Thus, it was possible, and quite useful, to try a set of quantitative data elicited from the doctor, compare the course of action recommended by the model to that recommended by the doctor, and to re-examine the probabilities and utilities involved when the two recommendations did not concur. Through this process of trial and error, it was soon easy to identify situations in which probabilities were not representative of the doctor's judgments or in which the utilities were not representative of his relative preferences. In fact, the first effort was to make small changes in utilities until the model worked, essentially tuning the model. The idea was to make it work, if possible, then

* See Gorry⁽²⁾ for a description of the program used.

re-evaluate the procedures for obtaining data and the structure of the problem to determine if they could be improved. This procedure not only increased intuitive understanding of the interaction of probabilities and utilities, but it revealed some of the limitations of the model being used.

The syndrome, acute renal failure, was selected for several reasons. First, the diagnosis of the cause of the acute renal failure can be a non-trivial problem. It is often necessary to try a treatment before a definite diagnosis has been confidently determined. Second, there are a number of tests available to a doctor which can be useful in the diagnosis, but with which are also associated varying degrees of risk, discomfort, and other sacrifices categorizable as test costs. Third, the diseases possibly responsible for the syndrome are very diverse in their seriousness and their responses to available treatments. In some cases, neither the best available treatment nor any of a number of mis-treatments is likely to significantly affect the course of the disease. On the other hand, there are situations in which correct diagnosis and treatment is likely to be very successful, but mistreatment can result in grave consequences. Last, but most important, acute renal failure was chosen because there were doctors available who were not only experienced in the diagnosis and treatment of the syndrome, but had expressed an interest in the project and a willingness to contribute their time and expertise.

This last consideration was particularly significant to the

accomplishment of this project, and it will be crucial to any attempts to further investigate the practicalities of a computer-aided diagnostic-treatment system. As will be illustrated in later portions of this paper, the difficulties encountered in gathering the necessary information were primarily those of communicating with the doctors. These problems in communication evolved from two sources. The vocabulary frequently invoked in medical practice and that used in quantitative analysis are different. Words and phrases necessary for the description of a disease or medical procedure are not familiar to the practice of quantitative analysis. Conversely, doctors are not conditioned to some of the expressions and mathematics common to quantitative analysis. So, it takes a great deal of time and conversation to develop a common vocabulary satisfactory to the problem being considered. The other source of communication difficulties was the experimental nature of the project. Not only was there uncertainty as to the appropriateness of applying utility theory to a diagnosis-treatment problem, but there was uncertainty as to the doctors' ability to express the judgments and preferences required. Questions that seemed explicit and appropriate to the analyst often were confusing or meaningless to the doctor. Similarly, some aspects of medical practice considered by the doctor in answering posed questions did not appear to be pertinent to the information being sought. Therefore, in addition to establishing a common vocabulary, it was necessary to fully discuss the implications of questions to both the mathematical model being considered and the medical

problem being described. There was also a necessity for determining which aspects of the medical problem should be included or specifically excluded in responding to questions. Thus, eliciting information from a doctor involved much more than simply presenting him with a questionnaire or a few direct questions. It required an intensive exchange of concepts and interpretations between the doctors and the analyst and proved to be demanding of the doctors' time and interest. It is because the doctors involved in this project were so cooperative and generous in the giving of their time and advice that the project was at all possible.

3.2 The Model

The model developed for the analysis of acute renal failure is diagrammed in Figure 3. The diseases considered, the tests and their possible results, the treatments, and the consequences are presented in Tables 1 through 4. Note that Figure 3 really represents an entire decision tree. A test need not be given before a treatment can be chosen (Part of the analysis is to determine whether or not to give a test at all.). Similarly, the model does not necessarily restrict the analysis to one stage of testing as shown in the figure. A sequence of any desired number of testing stages can be evaluated before choosing a course of action. So, the model is basically the same as that described earlier in Figure 2.

Several simplifying modifications, however, have been made. First, the test and treatment costs were considered to be independent of the actual disease as well as each other. Therefore, these costs are effectively predetermined tolls which must be paid at the time a test is given or treatment is carried out, rather than additional consequences of a testing-treatment policy. This modification greatly simplified data collection and actual computation since the decision tree beyond the treatment decision node remains the same regardless of testing accomplished. Likewise, the cost of any particular test or treatment was constant regardless of previous or subsequent testing.

In addition, it was assumed that in any specific case, the syndrome is actually a result of only one disease; the model does not allow the patient to be suffering from more than one disease at a time. Similarly,

CHOOSE ONE OF
SIX TESTS

OBTAIN ONE OF
EIGHTEEN
TEST RESULTS
TREATMENTS

ONE OF THE
FOURTEEN
STATES OF
NATURE EXISTS

OBTAIN ONE OF
THREE
CONSEQUENCES

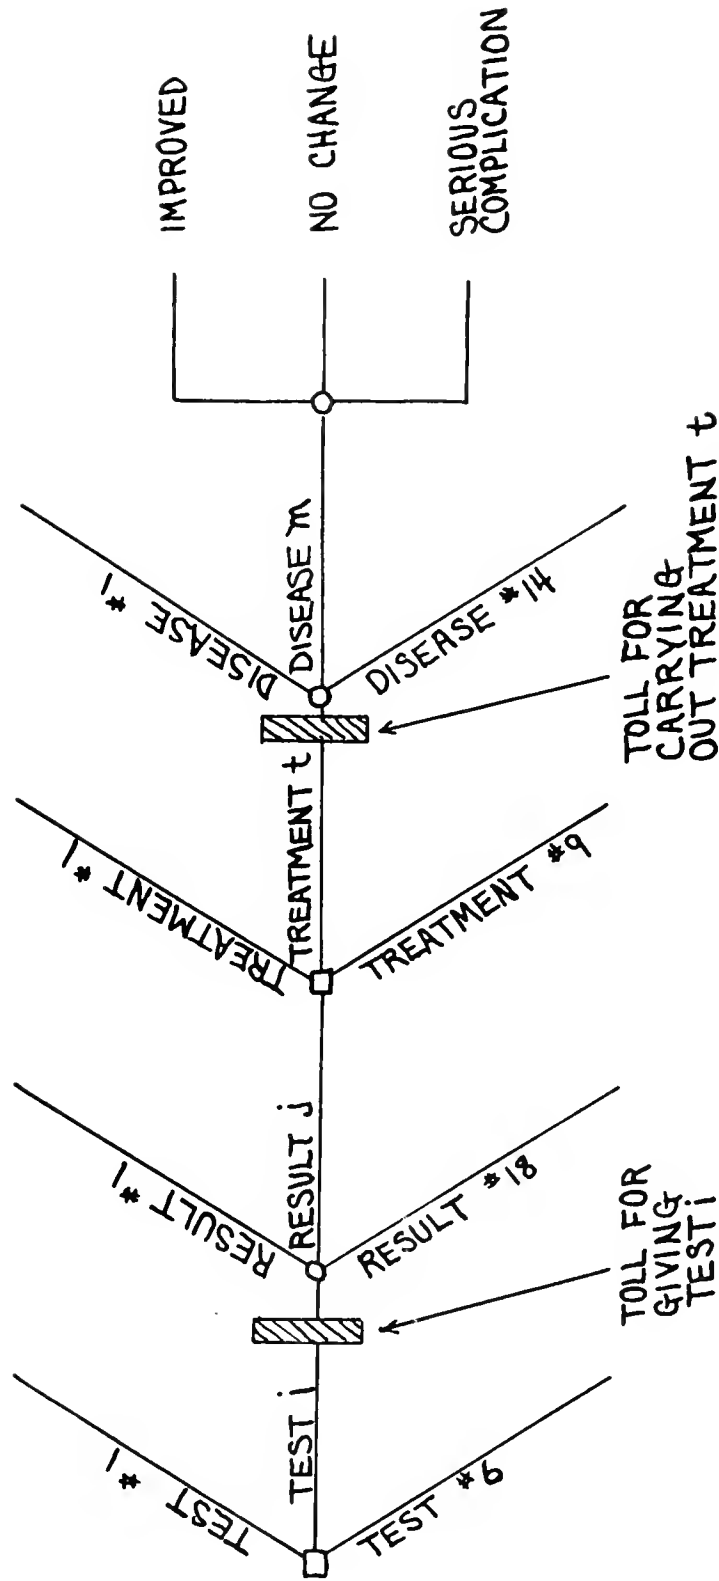


FIGURE 3

TABLE 1
DISEASES

	<u>DISEASES</u>	<u>ABBREVIATION</u>
D01	Acute Tubular Necrosis	ATN
D02	Functional Acute Renal Failure	FARF
D03	Obstruction of Urinary Tract	OBSTR
D04	Vasculitis	AG
D05	Cortical Necrosis	CN
D06	Hepatorenal Syndrome	HS
D07	Pyelonephritis	PYE
D08	Atheroembolism	AE
D09	Renal Infarction	RI
D10	Renal Vein Thrombosis	RVT
D11	Acute Interstitial Nephritis	AIN
D12	Scleroderma	SCL
D13	Acute Renal Disease/Chronic Disease	ARDCD
D14	Malignant Hypertension	MH

TABLE 2
TREATMENTS AND CONSEQUENCES

TREATMENTS

T01	No Therapy
T02	IV Fluids
T03	Mannitol
T04	Surgery for Obstruction
T05	Steroids
T06	Antibiotics
T07	Surgery for Clot
T08	Antihypertensive Drugs
T09	Heparin

CONSEQUENCES

R01	Patient's Condition Improved
R02	No Change in Patient's Condition
R03	Patient Experiences a Serious Complication

TABLE 3

TESTS

<u>TESTS</u>		<u>POSSIBLE RESULTS*</u>
T10	Biopsy	R04 thru R15
T11	Retrograde	R16, R17
T12	Plain Film of Abdomen	R18, R19
T13	Catheterization	R16, R17
T14	Arteriography	R20, R21

*See Table 4 for listing of test results.

TABLE 4
TEST RESULTS

TEST RESULTS

R04	No Specific Finding
R05	ATN
R06	AG
R07	CN
R08	HS
R09	PYE
R10	AE
R11	RI
R12	AIN
R13	SCL
R14	ARDCD
R15	MH
R16	OBSTR
R17	NO OBSTR
R18	Enlarged Bladder Visible
R19	Bladder Not Visible
R20	Blood Clot
R21	No Blood Clot

there is no provision for simultaneously carrying out more than one treatment. Both of these assumptions simply limit the size of the model without affecting its basic structure.

Further simplification was achieved by assuming that the utilities of the three possible consequences of any treatment not only reflected the patient's condition subsequent to the treatment, but summarized all future costs and consequences. Hence, in selecting a course of action, the model used the consequence of a treatment as the horizon for the analysis. In real life, of course, the problem might not terminate at that point. If a treatment results in other than improvement in the patient's condition, there is still a diagnosis-treatment decision confronting the doctor. To consider such a situation using the present model, the a priori probabilities of the diseases would have to be updated to reflect the experience of the ineffective treatment and the problem recycled, once again determining a testing-treatment policy.

One aspect of the costs associated with tests and treatments requires amplification. Costs can be conveniently divided into risk and non-risk categories. The category of risk is the probability that a serious consequence will develop pursuant to the administration of a specific test or treatment. The category of non-risk may include monetary costs, patient discomfort, time, hospital facilities used, or whatever other non-risk factors the doctors may wish to consider. It should be realized that the risk of testing is a consequence, not a test result; it does not provide information about the cause of the acute renal failure since the test costs are being considered independent of

the actual disease. Thus, the toll for performing a test must include both risk and non-risk associated categories of cost. The risk associated with a treatment, however, cannot be considered independent of the actual disease, but is, in fact, reflected in the conditional probabilities obtained from the doctor: his judgment of the likelihood of each consequence occurring given each combination of diseases and treatments. Thus, the tolls associated with treatments need only consider the non-risk category of costs.

3.3 The Probabilities

Because the emphasis of this project was on utility assessment, no effort was made to invoke special procedures for determining the required probabilities. All probabilities were strictly subjective estimates made by one of the doctors involved in the project.* There was no study of historical data. Nor was there an attempt to elicit distributions over the probabilities.

Dr. #1 was simply asked to directly estimate the required probabilities based on his own best judgment. If he was aware of statistical studies that had substantiated some of the data, he was cautioned to use the results of those studies only if he really believed them. The purpose of subjective probabilities was to quantify the doctor's own judgments, whether these judgments were accurate or not.

There are three types of probabilities required for implementation of the model: the a priori (initial) probabilities of the diseases, the conditional probabilities of the test results given the tests and diseases, and the conditional probabilities of the consequences given the treatments and diseases. The a priori probabilities are the inputs to the model and change according to the specific medical case that is being processed. They represent the sum total of all information gained about a patient prior to the implementation of the model; the

* Two doctors participated in the project. They will be referred to as Dr. # 1 and Dr. # 2.

data from the patient's medical history, results of routine tests, findings from a physical examination, and the likelihood of the possible diseases occurring are all reflected in the a priori probabilities.

The conditional probabilities are a semi-fixed part of the model. They can be changed if experience warrants a revision, but they normally form a data base which remains consistent from case to case. The conditional probabilities of the test results given the tests and diseases (Appendix 2) represent the likelihood that any patient who has a specific disease and is administered a specific test will demonstrate each of the possible test results. This information is indicative of the reliability of a test and must reflect the possibility of human error as well as test inadequacies. The conditional probabilities of consequences given the treatments and diseases (Appendix 1) represent the likelihood that any patient who has a specific disease and is administered a specific treatment will experience each of the possible consequences. It is a measure of the effectiveness of each of the treatments.

The actual eliciting of conditional probabilities from the doctor took only minutes. There were, however, several interesting aspects of this portion of the project. Dr. # 1, who provided all the subjective probabilities, had previous experience with both computers and their application to problems in medical diagnosis. He also had worked with subjective probabilities and had some intuitive understanding of Bayes' Theorem. He was aware that a zero probability expressed a judgment of impossibility and, therefore, eliminated certain options with no oppor-

tunity for recall later in the case. Thus, a zero probability could exclude from consideration one of the possible diseases. The result of this understanding was the estimation of many probabilities as .01. This meant that those test results or consequences being considered were effectively not possible except through accident or freak circumstances, but he wished to keep open his options for later in the problem.

The doctor's approach to assigning probabilities to his judgments was usually to identify the possible but unlikely occurrences, assign low probabilities to them, then allot the remaining probability to the most likely occurrence. Sometimes the unlikely occurrences were first considered as a group, a probability assigned to that group, then later divided among the individual events. This accounts for such subjective probabilities as .87, .29, .09, etc.

It was also discovered that in some situations the doctor felt he could not make the required judgment with any great degree of confidence. These situations arose primarily when considering the likelihood of consequences occurring as a result of mistreatment. The doctor's reaction was that if a patient had the disease specified, the particular treatment being considered was not appropriate and had never, to his knowledge, been attempted. Obviously, any doctor who did grossly mistreat a disease would not be anxious to acknowledge that he had done so, let alone document and publish the case. So, there were times the doctor was simply asked to guess. The rationalization behind this approach was that doctors, either consciously or unconsciously, make

similar judgments daily, and their guess should be at least as good as anyone elses.

3.4 The Utilities

Chapter 1.3 discusses utility theory and the justification for using the maximization of expected utility as the criteria for choosing among available courses of action. Discussed here will be the techniques used for constructing a utility scale. Later sections will report the experiences of actually eliciting preferences from the doctors and applying the information so gained to the model.

There are three categories of preferences over which utilities are to be assessed: the preferences for test costs, which include both the risk and non-risk aspects; the preferences for treatment costs, which include only the non-risk aspects; and the preferences for the consequences. Ultimately, these preferences must be quantitatively expressed on a common utility scale. Initially, however, the approach was to establish utility scales over each of the categories individually, then to seek a way of merging them into a common scale.

The technique for assessing utilities is derived directly from the third assumption of utility theory, repeated here for convenience:

Assumption 3. If prospect A is preferred to prospect B, which is in turn preferred to C, then there is a mixture of A and C which is preferred to B, and a mixture of A and C to which B is preferred.

The mixtures referred to in Assumption 3 are probabilistic in nature. That is, they are similar to lotteries; they represent a chance at attaining prospect A and the complementary chance at attaining prospect C. If the lottery is greatly biased towards prospect A

(meaning that there is a sufficiently good chance of attaining prospect A) an individual would prefer the uncertainty of the lottery to the certainty of prospect B, since the order of preference is A - B - C. On the other hand, if the lottery is sufficiently biased towards the least preferred prospect C, an individual would prefer the certainty of prospect B. Presumably, then, there is some lottery of prospects A and C such that an individual is indifferent in his preference between the certain prospect B and the lottery.

If the upper and lower limits of a utility scale are arbitrarily set and assigned to prospects A and C respectively, the utility of prospect B can be computed using the provisions of Property 2 of a utility scale, repeated below:

Property 2. If B is a prospect where, with probability p , the individual faces prospect A and with prospect $1-p$ he faces prospect C, then $\text{Utility of B} = p(\text{Utility of A}) + (1-p)(\text{Utility of C})$. Thus, a scale has been constructed which reflects the relative preferences of prospects A, B and C. Similarly, any prospect for which the preference falls between the extremes of prospects A and C can be represented on the scale by identifying the appropriate indifference lottery of prospects A and C.

The application of this technique of utility assessment to preferences relevant to the medical diagnosis-treatment problem primarily involves posing the appropriate lottery choices in such a manner that only the desired preferences are involved. For example, in trying to

assess utilities over preferences for treatments, it has been explained that only the non-risk aspects of treatment tests should be considered. A doctor, however, in weighing the merits of a proposed treatment, does not separate the treatment costs into risk and non-risk categories. Nor does he divorce preferences for the costs as a whole from other aspects of the case he is considering. He invokes his training, experience, and intuition to simultaneously evaluate all factors bearing on the specific case at hand. Thus, to assess utilities, it is necessary to pose questions which enable the doctor to mentally strip away all non-pertinent considerations and focus only on issues material to the preferences desired.

CHAPTER 4

IMPLEMENTATION OF THE MODEL

4.1 The Methodology

The first step in implementing the model was to first establish the data structure: that is, the conditional probabilities and the utilities. The doctor was next asked to fabricate hypothetical cases that could be used to check the performance of the model. These cases consisted of a priori probabilities over several diseases and the course of action the doctor thought he would choose in such a situation. (Table 5). The given a priori probabilities for each case were then introduced into the model, the appropriate calculations performed by the computer, and the model's recommended course of action compared to that suggested by the doctor. Thus, the conditional probabilities and the utilities were elicited prior to implementation of the model and became the data structure of the model. The inputs to the model, at the time of implementation, were the a priori probabilities associated with the hypothetical case being studied. The output was a recommended course of action and a ranking of decision alternatives.

Conditional probabilities were assigned once, in the manner described earlier. Although they were occasionally reviewed when the model's performance was questioned, only two were ever changed. In most cases, the doctor felt confident that the conditional probabilities initially provided were as close as he could come to quantitatively

TABLE 5
HYPOTHETICAL CASES

<u>CASE NUMBER</u>	<u>DISTRIBUTION</u>
C01	OBSTR .75 ATN .12 AG .13
C02	FARF .90 ATN .05 AG .05
C03	ATN .85 FARF .10 AG .05
C04	ATN .50 OBSTR .50
C05	MH .75 ARDCD .20 SCL .05
C06	AG .75 MH .25
C07	MH .50 AG .50
C08	RI .50 CN .50
C09	PYE .80 AG .20
C10	PYE .40 AG .40 AE .20
C11	FARF .50 ATN .50
C12	AG .77 ARDCD .15 MH .08
C13	PYE .60 ATN .30 OBSTR .10
C14	PYE .10 OBSTR .60 ATN .30
C15	HS .60 ATN .20 OBSTR .20
C16	RI .50 CN .20 AG .30
C17	SCL .50 MH .50
C18	RVT .50 ATN .30 FARF .20
C19	ATN .50 AE .50
C20	ATN .35 FARF .35 AE .30
C21	OBSTR .05 AG .50 AE .45
C22	SCL .10 ATN .25 OBSTR .65
C23	SCL .10 CN .30 AG .60
C24	SCL .50 CN .50
C25	RVT .20 SCL .20 ATN .60
C26	MH .50 PYE .40 RI .10
C27	FARF .30 ATN .50 OBSTR .20
C28	FARF .50 ATN .10 OBSTR .40

expressing his judgments. Therefore, when the model did not seem to perform as desired, it was the utilities and methods of assessing utilities that were examined and, usually, revised.

4.2 The First Attempt at Utility Assessment*

Initial efforts to assess utilities were not entirely successful. However, since these efforts did lead to a better understanding of both the model and the issues being confronted, the procedures and results will be described.

Utilities were assessed separately over each of the three categories of preferences using the technique discussed in Section 3.4. The consequences were considered first. To elicit the necessary preference relationships, the following question was posed:

Consequences. Would you prefer a testing-treatment policy that was sure to result in "no change" in the patient's condition, or would you prefer a policy which offered a 50% chance of the patient's condition being "improved" and a 50% chance of his experiencing a "serious complication"?

The question, as posed, was apparently explicit to the doctor. He understood that only preferences for the three consequences were involved, irrespective of the particular patient, the specific disease, or the testing-treatment policy. His choice was between a sure consequence and a lottery offering a chance at a very undesirable consequence and the complementary chance at a very desirable consequence.

He rejected the lottery. The chance at the undesirable consequence was then reduced until the doctor was indifferent between the conse-

* Initially, only Dr. # 1 was involved in the project. The participation of Dr. # 2 is described in section 4.4.

quence "no change" and the lottery. When that point was reached, arbitrary limits of a utility scale were assigned to the extreme consequences, and the utility of "no change" was computed. Results are displayed in Tables 6 and 7.

When considering the treatments, the doctor found it helpful to first order the treatments according to his preferences. He was then presented with the following question for each of the specific treatments being considered:

Treatments. If all treatments were guaranteed to yield the same consequences, regardless of the patient or the disease being treated, would you prefer to carry out the specific treatment being considered, or would you prefer to take a p chance at "surgery (clot)" (the least preferred treatment) and a $1-p$ chance at carrying out "no therapy" (the most preferred treatment)?

The question was posed in the above manner so that it would be clear to the doctor that neither the effectiveness nor the risk of a treatment was to be considered. The results of the questioning and the evolved scale are also presented in Tables 6 and 7.

Similar procedures were followed in eliciting the utilities over the possible tests. However, the question posed reflected the necessity for considering both risk and non-risk aspects of the test costs:

Tests. If all tests were guaranteed to yield the same information about a patient's disease, would you prefer to give the specific test being considered, or would you prefer to take a p chance at giving a

TABLE 6
INDIFFERENCE LOTTERIES
(Dr. # 1 - First Attempt)

<u>CERTAINTY OPTION</u>	<u>LOTTERY OPTION</u>	
	(p) PRIZE 1	(1-p) PRIZE 2
CONSEQUENCES:		
No Change	.02 Serious Complication	.98 Improved
TREATMENTS:		
Surgery (obstr)	.33 Surgery (clot)	.67 IV Fluids
Steroids	.33 Surgery (obstr)	.67 IV Fluids
Antihypertensive Drugs	.20 Surgery (obstr)	.80 IV Fluids
Heparin	.10 Surgery (obstr)	.90 IV Fluids
Antibiotics	.08 Surgery (obstr)	.92 IV Fluids
Mannitol	.05 Surgery (obstr)	.95 IV Fluids
TESTS:		
Arteriography	.80 Biopsy	.20 Plain Film
Retrograde	.67 Biopsy	.33 Plain Film
Catheterization	.25 Biopsy	.75 Plain Film

EXPLANATION: The doctor expressed indifference between the specified certainty 1 option and the corresponding lottery option offering a p chance at Prize 1 and a (1-p) chance at Prize 2.

TABLE 7

INDIVIDUAL UTILITY SCALES
(Dr. # 1 - First Attempt)

<u>CONSEQUENCES</u>	<u>UTILITY</u>
Serious Complication	100*
No Change	2
Improved	0**
<u>TREATMENTS</u>	
Surgery (clot)	300*
Surgery (obstr)	100
Steroids	30
Antihypertensive Drugs	20
Heparin	10
Antibiotics	8
Mannitol	5
IV Fluids	0
No Therapy	0**
<u>TESTS</u>	
Biopsy	100*
Arteriography	80
Retrograde	66
Catheterization	25
Plain Film	0**

* Arbitrarily assigned as upper limit of the scale.

** Arbitrarily assigned as lower limit of the scale.

"biopsy" (the least preferred test) and a 1-p chance at giving a "plain film" (the most preferred test)?

The indifference lotteries and the resulting utility scale over the tests are again presented in Tables 6 and 7.

The next task was to merge the three utility scales into a common one. Exactly how to do this was not pre-determined. In fact, there was some doubt as to whether it could be done at all. As it turned out, the doctor had no trouble at all comparing tests and treatments, even though test costs included risk and treatment costs did not. Apparently, since both tests and treatments involve doing something to a patient, the factors considered in evaluating costs were very similar. The doctor was able to simply identify tests and treatments for which his preferences were about equal. In this manner, three tests and treatment scales were merged.

Merging consequences into the common scale, however, proved to be a much different chore. The doctor did not know how to compare a consequence with a treatment. They seemed to be incompatible; there was no apparent common ground for comparison. Certainly, the consequence of "improved" was preferable to all tests and treatments, and all tests and treatments were preferable to a "serious complication," but there seemed to be no way, either by direct comparison or the lottery technique, of relating the preferences for the consequences to the common utility scale for tests and treatments. Finally, the doctor decided his preference for "no change" was about equal to that for "anti -

hypertensive drugs," and the lower end of the treatment scale should be the same as "improved." The two scales were then stuck together using "no change" and "improved" as reference points. Needless to say, there was not much confidence expressed in the utility scale as it finally emerged (Table 8). Nevertheless, it was decided to enter the utilities into the data structure of the model, and to try a few cases to see what would happen.

As was expected, the performance of the model was not satisfactory. As the doctor commented at the time, it seemed to be simulating a medical student rather than a doctor. The model's recommended course of action for almost every case was the administration of the treatment "heparin." Looking at the conditional probabilities associated with heparin (Appendix 1), it is apparent that heparin is not particularly effective in treating any of the diseases, but neither is it very dangerous. The use of "heparin" in every case is a very cautious policy most likely to produce "no change" in the patient's condition. Since this is obviously not the attitude a doctor must take, the utilities of the three consequences apparently did not express the proper relationships among preferences.

A close examination of the model revealed at least partial insight into the reasons for the observations. Because the model being used considers a sequence of tests and treatments through only one treatment, the problem is being truncated at a point at which it may not actually terminate. If the outcome of a testing-treatment policy in a real

TABLE 8
UTILITY SCALE
(Dr. # 1 - First Attempt)

	<u>Utility</u>
Serious Complication	1000
Surgery (Clot)	300
Surgery (Obstr)	100
Biopsy	60
Arteriography	50
Retrograde	40
Steroids	30
No Change	20
Antihypertensive Drugs	20
Catheterization	15
Heparin	10
Antibiotics	8
Mannitol	5
IV Fluids	0
Plain Film of Abdomen	0
Improved	0

medical case were, in fact, "no change," the doctor would still have a diagnosis-treatment problem facing him. Therefore, in addition to the seriousness of the situation and the opportunity for further treatment, the utility of "no change" should include the expected utility of all subsequent courses of action. This meant that the utility of "no change" should be substantially larger than initially considered, and may even be greater than that of any of the tests and treatments.

The above conclusions were substantiated when the model was re-tested using a utility for "no change" in the 600 to 800 range. As long as the a priori probabilities indicated that the diagnosis was fairly certain, say 85% or higher, the model chose treatments that the doctor felt were reasonable. However, when the certainty of a diagnosis was lowered, so that two or more diseases seemed approximately equally likely to be the cause of the syndrome, the model consistently failed to recommend the use of tests to improve a diagnosis before treatment. Although the treatments recommended were appropriate for at least one of the likely diseases, the doctor felt that in many cases it would be desirable to improve diagnosis before treatment.

In reviewing the utility scale, the doctor also made the following observation. Although the scale derived might reflect his true preferences for test and treatment costs, when considering alternative courses of action in a real medical case, he did not visualize these costs as being individually distributed along a scale. Instead, he visualized them in groups, some having little or no cost, others having

moderate costs, and still others having very high costs. Within a group, he would choose a test or treatment solely on its usefulness in improving the patient's condition, no distinction being made in costs. Therefore in keeping with the project's goal of duplicating a doctor's decisions, it may be more representative of the doctor's attitudes to group tests and treatments, then assess utilities over the several groups, rather than consider preferences for individual costs.

4.3 The Second Attempt at Utility Assessment

Rather than continue to empirically adjust numbers, the utilities previously derived were abandoned, and the lessons learned from the first attempt were applied to improve the techniques of utility assessment. Tests and treatments were grouped according to their non-risk costs and a utility scale then established over the groups using the lottery technique (Table 9, column a). To account for the risk attributable to each test, the doctor was asked to estimate the probability that each test, independent of the disease or treatment, would lead to a "serious complication." The utility of each test was then the assessed utility of its non-risk costs plus some fraction of the utility of a "serious complication." (Table 9, columns b and c)

No attempt was made to assess utilities over the three consequences. Experience had indicated that the utility of "no change" would have to be about three quarters that of "serious complication" and that the utilities of all three consequences might have to be modified by trial and error. Therefore, utilities of 10,000, 7,500, and 0 were assigned to "serious complication," "no change," and "improved" respectively.

This still left unresolved the problem of how to relate the utilities of the consequences to the utilities of the tests and treatments. Recall that during the first attempt, these two scales were just "stuck" together — admittedly not a very satisfying solution to the problem. This time, the following approach was tried. The doctor was asked to make a choice between two courses of action. One course of

TABLE 9
UTILITIES
(Second Attempt)

<u>Group</u>	<u>Tests and Treatments</u>	a. <u>Non-risk Utility</u>	b. <u>Risk*</u>	c. <u>Utility</u>
I	Surgery (Clot)	1000		1000
II	Surgery (Obstr)	300		300
III	Biopsy	100	5%	600
	Arteriography	100	1%	200
IV	Retrograde	50	25%	300
	Steroids	50		50
	Antihypertensive Drugs	50		50
V	Catheterization	1	2%	200
	Heparin	1		1
	Antibiotics	1		1
	Mannitol	1		1
	IV Fluids	1		1
	Plain Film	1		1
	No Therapy	1		1

* Probability that a test independent of the disease or treatment, will lead to a serious complication.

<u>Consequences</u>	d. <u>Initial Utility</u>	e. <u>Adjusted Utility</u>
Serious Complication	10,000	6,000
No Change	7,500	4,500
Improved	0	0

action offered an opportunity to carry out "surgery (clot)" with the knowledge that the treatment would result in the consequence "improved." In other words, there was no risk associated with the surgery; only the non-risk costs would be incurred. The alternative to surgery was to accept a p chance at "serious complication" and the complementary chance, $1-p$, at "improved." The p was then adjusted until the doctor expressed indifference between the two courses of action. The result, p equal to .10, was interpreted to mean that the utility of "surgery (clot)" should be about one tenth the utility of "serious complication." Although the same procedure was applied to other treatments in an attempt to establish several reference points for merging the two utility scales, all values of p were so close to zero that no useful differentiation was obtained. Accordingly, the lower extremes of both scales were considered to have negligible utility. The resulting utility scale over tests, treatments, and consequences is reported in Table 9, columns c and d.

The model's performance on ten of the hypothetical cases, using the established utilities, is presented in Table 10. The most obvious difference between the model's decisions and those of the doctor is the model's repeated selection of the test "biopsy" in situations in which the doctor would forego testing altogether in favor of immediate treatment. As can be appreciated by examining the conditional probabilities associated with "biopsy," a biopsy is a very definitive test; in most cases it results in an almost sure diagnosis with little chance of

TABLE 10
DECISIONS ON HYPOTHETICAL CASES
(Dr. # 1 - Initial Utilities)

<u>CASE</u>	<u>Dr. # 1</u>	<u>MODEL</u>
C01	T11	T10
C02	T02	T10
C03	T01	T10
C04	T11	T10
C05	T08	T08
C06	T05	T10
C07	T10	T10
C08	T14	T14
C09	T06	T10
C10	T10	T10

NOTE: The utilities used in the model were those initially assessed by Dr. # 1 during the second attempt at utility assessment.

misdiagnosis. Therefore, the model, by extensive use of the "biopsy," was being very cautious with regards to its treatments.

There are two factors which could produce this bias in the model. One is that the assessed utility of "biopsy" reflects lower costs or risk than the doctor perceives when he actually considers a case. In other words, the model's decisions might be more in accord with those of the doctor if the utility of "biopsy" were substantially raised. The second possible cause of the bias is that the assessed utility of a "serious consequence" might be much higher than that apparent from the doctor's decisions. So that lowering the utility of a "serious consequence" might improve the model's performance.

Because there seemed to be less confidence in the methods used to relate tests and treatments to consequences than in the assessment of utilities among the tests and treatments, the utility of a "serious consequence" (also that of "no change") was lowered about one third. The adjusted utilities for the consequences are displayed in Table 9, column e.

Using the adjusted utilities for the consequences, the model was again tested on the hypothetical cases. This time, the model's selected courses of action initially agreed with those of the doctor on twenty-two of the twenty-eight cases. However, when the doctor reviewed the six cases in which there was disagreement, he changed one of his decisions to agree with the model, indicating that the model's response was really a better decision. And for two of the cases, he stated that the

model's choices were just as good as his, both being appropriate courses of action in these situations. So, there were essentially only three of the twenty-eight cases in which the model did not succeed in duplicating the doctor's decisions. (Table 11)

Notice also that neither "catheterization" nor "plain film," both tests for an "obstruction," are ever recommended by either the doctor or the model; whereas the "retrograde," which has a higher utility (Table 9), is chosen in several cases. Use of "catheterization" is restricted because it can only test for an obstruction in the neck of the bladder; the retrograde, on the other hand, can detect an obstruction in either the neck of the bladder or in the upper receiving tract between the bladder and the kidney. So when "obstruction" is allowed to refer to either of these areas, the "retrograde" is a much more reliable test, and that reliability is reflected in the conditioned probabilities (Appendix 2). The model can easily be modified to reflect a more realistic role for "catheterization" by dividing the disease "obstruction" into two new diseases which specify the general location of the obstruction. Of course, new probabilities would then have to be estimated to be in accordance with the new definitions.

The "plain film" is a different matter. In practice, a "plain film" of the abdomen would be followed by either "retrograde" or "catheterization" regardless of the results. So the doctors, themselves, appreciate that a "plain film" provides them with very little information. And because its results do not affect the decision

TABLE 11

DECISIONS ON THE HYPOTHETICAL CASES
(Dr. # 1 - Adjusted Utilities)

<u>CASE</u>	<u>Dr. # 1</u>	<u>MODEL</u>
C01	T11	T11
C02	T02	T02
C03	T01	T01
C04	T11	T11
C05	T08	T08
C06	T05	T05
C07**	T10	T08
C08***	T14	T09
C09	T06	T06
C10	T10	T10
C11*	T02	T10
C12	T05	T05
C13	T06	T06
C14	T11	T11
C15	T11	T11
C16	T14	T14
C17	T08	T08
C18***	T09	T02
C19	T01	T01
C20*	T02	T10
C21	T12	T12
C22	T11	T11
C23	T05	T05
C24	T09	T09
C25	T01	T01
C26*	T06 and T08	T10
C27	T11	T11
C28	T11	T11

* The model and Dr. # 1 disagree.

** The doctor later changed his decision to agree with the model.

*** The doctor stated that either his or the model's decision was appropriate.

NOTE: The utilities used in the model were those assessed by Dr. # 1 during the second attempt at utility assessment. The utilities for the consequences have been adjusted on the basis of past experience with the model.

regarding the next course of action, the obvious question is "Why use a plain film at all?". After pondering this question, the doctor concluded that it really is a useless test, and had this analysis been made before the model was developed, "plain film" would not be one of the tests considered.

4.4 Utility Assessment by Dr. # 2.

Up to this point in the project, the same doctor had helped develop the model, had provided all the sample cases and probabilities, and had participated in the utility assessment exercise. Further, as noted earlier, he had had considerable experience in working with probabilities, Bayes Theorem, and computers. So he had a reasonable understanding of the entire project and the procedures used for utility assessment.

The next step in the project was to try the model using utility assessments made by another doctor who was less familiar with the model and uninvolved in developing the utility assessment techniques. The primary reason for this was that the questions used for assessing utilities were formed through discussions with the doctor. The model required specific information, and it was necessary to exchange ideas and vocabulary to establish questions which asked for only that information. As a result, it was not certain whether the questions themselves, independent of the discussions, were adequate for eliciting the desired information.

Dr. # 2 was also familiar with various applications of computers in medicine, but he was not familiar with either the model being used or with utility theory. Nor had he ever been involved in any utility assessment exercises. So even though he was aware of the project and its objectives, he was free of any influences that the discussions, which led to development of the model and the questions, might have on

the actual utility assessment. It should also be noted that Dr. # 2 did not provide estimates of the conditional probabilities; he only assessed utilities. The two doctors have practiced medicine in the same clinic for about ten years. In fact, Dr. # 1, who helped develop the model and questions used in the project, studied medicine under Dr. # 2. So it would not be surprising if the two doctors expressed similar preferences and judgments.

To introduce Dr. # 2 to the project, he was given an explanation of both the model and utility theory to about the extent presented in Chapters 1 and 3 in this paper. He was then led through the utility assessment in the same manner as had been employed in the second attempt previously described. The resulting utility scale is shown in Table 12.

An interesting observation concerning non-risk costs emerged from the utility assessment exercise. In past attempts at utility assessment, there was never any attempt to determine exactly what factors the doctor considered in evaluating and comparing non-risk costs of a test or treatment. Examples were given so the doctor would understand the differentiation between risk and non-risk costs, and also so he would realize that cost could include many factors other than monetary expenses. But he was never told specifically which factors to consider, nor was he asked to state which factors he had considered. His instructions were to consider whatever he thought he would consider if he were actually caring for a patient. This time, the doctor declared emphatically that the only non-risk factor he felt actually affected his

-

decisions in practice was patient discomfort. Neither money, doctor's time, nor hospital facilities was ever considered a limitation or criteria affecting decisions concerning a patient suffering from acute renal failure.

There were only several minor problems encountered during the utility assessment. These all dealt with definition of the heretofore unspecified patient and some of the treatments. For example, the "retrograde," a test, is much less unpleasant for a female patient than for a male. Similarly, the oral administration of "antibiotics," a treatment, is less unpleasant than administration by a series of injections. So these factors might well affect utility assessment. However, as long as the doctor is consistent in considering the same factors when he evaluates the sample cases as he does when assessing utilities, then problems in definition should not affect the model.

Once again, however, there was difficulty relating the utilities of the tests and treatments to those of the consequences. The doctor was presented with the same choice offered during the last attempt — a choice between incurring the non-risk costs of a successful "surgery" and incurring the uncertain consequences of a lottery which has "improved" and "serious complication" as its prizes. Even though the surgery was considered very unpleasant for the patient when compared to other tests and treatments, when compared to a "serious complication" it could be almost trivial. In fact, the doctor thought he would accept no more than about a 1% chance at a "serious complication" in

lieu of the risk free surgery. Though this reply would indicate that the utility of a "serious complication" should be about one hundred times as large as that of "surgery," it was obvious from past experience that this figure far exceeded numbers for which the model might yield acceptable results. So, as was done previously, the utilities of the consequences were empirically set as indicated in Table 12.

Further limitations of the model were revealed when the doctor was asked to consider the hypothetical cases and to state what his courses of action would be in such situations. Apparently, the patient's condition could greatly influence the doctor's decisions. If the kidneys have not yet been damaged, the doctor can prescribe "no therapy" for the time being, or he can carry out treatment with one of the low cost drugs even though that treatment may not be the most likely cure. In other words, time is not critical. On the other hand, if the condition of the kidneys is deteriorating, the doctor must immediately determine the most likely disease, and carry out treatment directed specifically at that disease. In terms of the model, this means that the conditional probabilities of the consequences given a treatment and a disease may vary according to the patient's condition. Similarly, the utility of the consequence "no change" may depend on the patient's condition. Unfortunately, the model, as used, did not provide for different patient conditions. Nor was the patient's condition specified. So, when the doctor considered the sample cases, he had to assume a patient condition; and his assumption may or may not coincide with the assumptions made

TABLE 12
UTILITY SCALE
(Dr. # 2)

<u>Group</u>	<u>Test or Treatment</u>	<u>Non-Risk Utility</u>	<u>Risk</u>	<u>Utility</u>
I	Surgery (Clot	1000		1000
	Surgery (Obstr)	1000		1000
II	Biopsy	50	4%	450
	Retrograde	50	2%	250
III	Steroids	40		40
	Antihypertensive Drugs	40		40
IV	Arteriography	30	3%	330
	Heparin	30		30
	Catheterization	30	1%	130
V	Mannitol	1		1
	Antibiotics	1		1
	IV Fluids	1		1
	Plain Film	1		1
	No Therapy	1		1

Consequences

Utilities

Serious Complication

8000

No Change

5000

Improved

0

when the probabilities were estimated.

Despite these problems in definition, the model succeeded in duplicating the doctor's decisions on twenty-four of the twenty-eight hypothetical cases. Furthermore, in two of the four cases in which there was disagreement, the cause of the disagreement was traced to the conditional probabilities. (Remember that Dr. # 2 provided just the utilities, not the probabilities.) In both of these cases, when the probabilities were changed to reflect Dr. # 2's judgments, the model then duplicated his decisions. So, agreement between the model and Dr. # 2 was obtained on all but two of the twenty-eight cases.

It is also interesting to observe by comparing Tables 11 and 13 that even though Dr. # 1 and Dr. # 2 disagreed on the appropriate course of action for cases C06, C13, and C16, the model followed the appropriate doctor. When the utilities assessed by Dr. # 1 were used in the model, the model's decisions agreed with him. When the utilities were changed to those assessed by Dr. # 2, the model's decisions also changed to agree with Dr. # 2.

TABLE 13
DECISIONS ON THE HYPOTHETICAL CASES
(Dr. # 2)

<u>CASE</u>	<u>Dr. # 2</u>	<u>MODEL</u>
C01	T11	T11
C02	T02	T02
C03	T01	T01
C04	T11	T11
C05	T08	T08
C06	T08	T08
C07	T08	T08
C08	T09	T09
C09	T06	T06
C10	T10	T10
C11**	T02	T02
C12	T05	T05
C13**	T11	T11
C14	T11	T11
C15	T11	T11
C16	T09	T09
C17	T08	T08
C18	T09	T09
C19	T01	T01
C20	T10	T10
C21*	T11	T12
C22	T11	T11
C23	T05	T05
C24	T09	T09
C25	T01	T01
C26*	T06 & T08	T10
C27	T11	T11
C28	T11	T11

* The model and Dr. # 2 disagree.

** The model and Dr. # 2 agree only after the conditional probabilities were changed to represent this doctor's judgments.

NOTE: The utilities used in the model were those assessed by Dr. # 2. Except as noted above by **, the conditional probabilities were estimated by Dr. # 1.

CHAPTER 5

DISCUSSION

Previous chapters have explained the formulation of a utility theory model of the diagnosis and treatment of acute renal failure and have presented an account of attempts to implement that model. This chapter brings together several comments and assessments concerning the model and techniques implemented in this project and suggests areas for further research.

Perhaps the best place to begin an evaluation is to recall the objectives and the results. This project sought to use utility theory to develop a quantitative model, which was adaptable to use on a computer, and which could be used to predict or duplicate decisions made by a doctor in diagnostic-treatment situations. Because previous studies had indicated that the probabilistic aspects of the theory seemed appropriate for this type of problem, most of this project concentrated on using utility theory to find numbers, or utilities, which appropriately scaled a doctor's relative preferences for tests, treatments, and consequences.

To determine if the model, hence the utilities, were attaining that objective, the model's recommended courses of action on a number of hypothetical medical cases were compared to doctors' responses to those same cases. Had there been very little or no agreement between the model and the doctors, one would have had to conclude that at least

the techniques used to implement the model were inadequate. There might also have been some grave reservations about the structure of the model.

However, the model's performance on the hypothetical cases corresponded well to the doctors' analyses of the cases. With one doctor there was agreement on twenty-five of the twenty-eight cases. With the other, once probabilities were modified to his satisfaction, there was agreement on twenty-six of the cases. Furthermore, when the doctors disagreed on a case, the model duplicated each doctor's decision when using his utilities. Certainly such success has not only been rewarding, but has given reason to be optimistic about future applications of utility theory to diagnostic-treatment problems.

Of course, it would be premature to conclude that all problems associated with applying utility theory to medical diagnosis and treatment have been solved. There were, admittedly, limitations to the model used in this project. Many factors which enter into a doctor's decision — such as patient identification, patient condition, and duration of treatment — were not included. Some were purposely excluded to simplify the model; the exclusion of others became apparent only as the project progressed. The result of such limitations, beside occasional confusion, was that the model's performance could not be tested on actual medical cases; hypothetical cases had to be used instead. Although this was satisfactory for the goals of this project, it does limit the significance of the model's success.

On the other hand, none of the model's limitations seemed to be inherent in the characterization of medical diagnosis as decision making under risk. On the contrary, the doctors felt that this is a good way of approaching medical decisions. After all, the procedure of consciously identifying various courses of action, their known costs, their risks, and their possible consequences is simply a very methodical way of accomplishing the same type of analysis that doctors invoke intuitively. The primary difference is that doctors do not try to quantify either their judgments or preferences. Yet, the doctors who participated in this project were able, with only several exceptions, to readily provide the information necessary to make those quantifications.

The most obvious next step in exploring the use of utility theory in medical diagnosis and treatment is to expand the present model to the extent that it can be tested against actual medical cases. Those limitations that have been mentioned throughout this paper can all be easily overcome by further definition of either the patient, the treatments, or the diseases. Certainly, the model will be larger and more complex: more diseases, more treatments, and more probabilities. But the basic model and a good computer program already exist. The only problem might be collecting appropriate medical cases from a clinic.

There is also potential use for this type of model in the field of medical education. A computer program might be developed which could require a student to make a decision on a hypothetical medical case,

and could then critique the student's decision, pointing out such factors as effectiveness of a test or treatment, risk involved, or patient discomfort expected.

The use of directly estimated subjective probabilities is another area in which further research might be helpful. There was no effort in this project to confront the many issues relevant to probability estimation or the use of subjective probabilities. The doctor merely stated his estimates, and they were thereafter faithfully used in the model. It was pointed out, though, that one doctor had previous experience in estimating probabilities. So, there was some prior indication that this approach would prove adequate, and the results seem to substantiate that assumption.

A close look at the overall effect of the probabilities will suggest that this casual approach of direct estimation may be a good technique. Even for this limited project it would have taken a major effort, over a long period of time, to statistically determine the required probabilities with any high degree of confidence. On the other hand, direct estimation took only a few minutes. And even though the doctor's estimates may not be individually accurate, he usually knows whether a probability should be high, low, or somewhere in the middle. As long as the numbers he chooses reflect his judgment about the relative effectiveness of the tests or treatments on each of the diseases, the fact that he has difficulty distinguishing between a 0.20 and a 0.25 does not significantly affect the model. As a practical

matter, should such small variations in the probabilities result in a change in the recommended test or treatment, a doctor would probably agree that either decision is appropriate for the case.

The situations in which direct estimation seemed least satisfying were those concerning rare, unusual or unheard of mistreatments. The doctor simply did not know what would happen if a patient was so grossly mistreated; so his probability estimates were just guesses. Even these situations had their contribution to the project: they helped identify areas in which medical knowledge needs to be improved. Experiments or statistical studies in these few areas may be all that is necessary to raise a doctor's confidence in some of his guesses and provide entirely adequate probabilities for utility theory type models.

The direct use of utility theory to establish scales over tests, treatments and consequences also had its successes and failures. In scaling the non-risk costs of tests and treatments, the lottery technique was very satisfactory. The doctors had no trouble responding to the questions and were satisfied with the resulting scales. In fact, the practice of grouping several tests and treatments having similar non-risk costs proved quite appropriate to the type analysis a doctor makes. The key issue in these situations was whether, within a group, the non-risk cost would affect a doctor's decision. Even though there may be differences in non-risk costs among the tests or treatments in a group, and even though those differences may be very apparent when isolated in a lottery choice situation, as a practical matter, the risk

or effectiveness of the tests or treatments so greatly overshadows their differences in non-risk costs that those non-risk costs never really affect a decision. In fact, the risks associated with the various tests considered in this project sufficiently overshadowed the non-risk costs that those costs might safely be ignored. More experience in dealing with medical decisions might lead to other methods for simplifying utility assessment.

One of the doctors suggested that patient discomfort is the only non-risk cost that needs to be considered when assessing utilities over tests and treatments. If this is a general attitude among doctors, it would be interesting to determine if the doctors and their patients agree on preferences for various forms of discomfort, or if there is a way of either qualitatively or quantitatively defining categories of discomfort.

Attempts to apply the lottery technique to the three consequences and to then relate those consequences to the tests and treatments proved frustrating. This project never did establish a satisfactory means of quantifying those relationships. The stumbling block was the extreme disparity between the undesirability of the "serious complications" and that of the non-risk costs of the treatments. Isolating the "serious complications" from a real medical situation and focusing on their undesirability seemed to magnify their gravity. When efforts were made to quantify this severe attitude, the effect was the establishment of a very cautious model, one which would recommend treatment only after a very sure diagnosis had been obtained.

The trade-off operating here is that the higher the utility of a "serious complication," the more the probabilities, hence the risk, dominate the model. As the utility of "serious complication" is lowered, the more the relative non-risk costs influence the decision. In this project, both the computer program used and the on-line mode of operating the computer made it easy to observe this trade-off. Utilities could be easily changed and a case re-run on the model. Immediate response from the computer allowed convenient comparison of present and past performance of the model and evaluation of the affect of the changed utilities. It also increased intuitive understanding of the model and enabled the users to identify means of improving the model's performance.

Because there were only three consequences involved, this man-machine interaction was sufficient for adjusting consequence utilities until the model seemed to work. In more complex models, however, it may not be possible to adequately represent all consequences in so few categories, and trial and error methods might prove inadequate. A useful approach might be to define a category of consequences that is comparable to treatments, and use it both to establish a meaningful scale over the consequences and to quantify the relationship between consequences and treatments. In any event, this comparison is considered to be the least satisfying aspect of this project. Any improvement in the techniques used to make the comparison would greatly enhance the appeal of using utility theory to model medical diagnosis-treatment problems.

BIBLIOGRAPHY

1. Chernoff, Herman and Lincoln E. Moses. Elementary Decision Theory. New York: John Wiley and Sons, Inc., 1959.
2. Gorry, G. Anthony. "Strategies for Computer-Aided Diagnosis," Mathematical Biosciences, 2, 293-318 (1968).
3. Gorry, G. Anthony, "A System for Computer-Aided Diagnosis," Project MAC Technical Report MAC-44 (Thesis). Cambridge, Massachusetts: Massachusetts Institute of Technology, 1967.
4. Gorry, G. Anthony, and G. Octo Barnett. "Sequential Diagnosis by Computer," The Journal of The American Medical Association, 205: 849-854 (September, 1968).
5. Lodwick, Gwilym S., and others. "Computer Diagnosis of Primary Bone Tumors: A Preliminary Report," Radiology, 80: 273-275 (February, 1963).
6. Luce, Duncan R., and Howard Raiffa. Games and Decisions: Introduction and Critical Survey. New York: John Wiley and Sons, Inc., 1968.
7. Overall, John E. and Clyde H. Williams. "Conditional Probability Program for Diagnosis of Thyroid Function," The Journal of the American Medical Association, 183, 5: 95-101 (February, 1963).
8. Raiffa, Howard. Decision Analysis: Introductory Lectures on Making Decisions Under Uncertainty. Reading, Massachusetts: Addison-Wesley Publishing Company, 1968.
9. Rubel, Ronald R. "Decision Analysis and Medical Diagnosis and Treatment," unpublished doctoral thesis, Graduate School of Business Administration, Harvard University, Cambridge, Massachusetts, 1967.
10. Warner, Homer R., Alan F. Toronto, and L. George Veasy. "Experience with Bayes Theorem for Computer Diagnosis of Congenital Heart Disease," Annals of the New York Academy of Sciences, 115, 2: 559-566 (1964).

APPENDIX 1

CONDITIONAL PROBABILITIES OF CONSEQUENCES

EXAMPLE

T03

D05 .00 .80 .20

Explanation:

If the cause of the acute renal failure is known to be disease D05 and the treatment T03 is carried out, the probability that consequence R01 will result is .00, the probability that consequence R02 will result is .80, and the probability that consequence R03 will result is .20.

T01

D01	.20	.75	.05	D08	.05	.70	.25
D02	.05	.45	.50	D09	.05	.05	.90
D03	.05	.85	.10	D10	.10	.75	.15
D04	.20	.70	.10	D11	.20	.75	.05
D05	.10	.80	.10	D12	.05	.85	.10
D06	.05	.45	.50	D13	.20	.75	.05
D07	.05	.70	.25	D14	.05	.75	.20

T02

D01	.05	.35	.60	D08	.01	.39	.60
D02	.90	.01	.09	D09	.01	.39	.60
D03	.01	.39	.60	D10	.20	.20	.60
D04	.05	.35	.60	D11	.01	.39	.60
D05	.05	.35	.60	D12	.01	.39	.60
D06	.05	.25	.70	D13	.01	.39	.60
D07	.05	.35	.60	D14	.01	.19	.80

T03

D01	.20	.70	.10	D08	.01	.90	.09
D02	.01	.39	.60	D09	.01	.90	.09
D03	.01	.49	.50	D10	.01	.90	.09
D04	.01	.90	.09	D11	.01	.90	.09
D05	.01	.90	.09	D12	.01	.90	.09
D06	.05	.76	.19	D13	.01	.90	.09
D07	.01	.90	.09	D14	.01	.90	.09

T04

D01	.00	.80	.20	D08	.00	.50	.50
D02	.00	.50	.50	D09	.00	.50	.50
D03	.80	.05	.15	D10	.00	.80	.20
D04	.00	.80	.20	D11	.00	.80	.20
D05	.00	.80	.20	D12	.00	.80	.20
D06	.00	.80	.20	D13	.00	.80	.20
D07	.01	.79	.20	D14	.00	.50	.50

T05

D01	.01	.80	.19	D08	.01	.80	.19
D02	.01	.80	.19	D09	.01	.80	.19
D03	.01	.80	.19	D10	.01	.80	.19
D04	.50	.31	.19	D11	.50	.31	.19
D05	.10	.71	.19	D12	.10	.60	.30
D06	.05	.70	.25	D13	.20	.61	.19
D07	.05	.76	.19	D14	.01	.59	.40

T06

D01	.01	.90	.09	D08	.01	.90	.09
D02	.01	.90	.09	D09	.01	.90	.09
D03	.01	.90	.09	D10	.01	.90	.09
D04	.01	.90	.09	D11	.01	.90	.09
D05	.01	.90	.09	D12	.01	.90	.09
D06	.01	.90	.09	D13	.01	.90	.09
D07	.86	.05	.09	D14	.01	.90	.09

T07

D01	.00	.60	.40	D08	.05	.50	.45
D02	.00	.30	.70	D09	.50	.10	.40
D03	.00	.60	.40	D10	.30	.30	.40
D04	.00	.60	.40	D11	.00	.60	.40
D05	.00	.30	.70	D12	.00	.60	.40
D06	.00	.10	.90	D13	.00	.60	.40
D07	.00	.40	.60	D14	.00	.60	.40

T08

D01	.01	.90	.09	D08	.05	.87	.08
D02	.01	.29	.70	D09	.01	.90	.09
D03	.01	.85	.14	D10	.01	.90	.09
D04	.30	.60	.10	D11	.01	.90	.09
D05	.01	.29	.70	D12	.01	.40	.59
D06	.01	.29	.70	D13	.35	.56	.09
D07	.05	.85	.10	D14	.40	.51	.09

T09

D01	.01	.94	.05	D08	.01	.89	.10
D02	.01	.94	.05	D09	.10	.85	.05
D03	.01	.94	.05	D10	.20	.75	.05
D04	.01	.94	.05	D11	.01	.94	.05
D05	.40	.40	.20	D12	.01	.94	.05
D06	.01	.59	.40	D13	.01	.94	.05
D07	.01	.94	.05	D14	.01	.94	.05

APPENDIX 2

CONDITIONAL PROBABILITIES OF TEST RESULTS

EXAMPLE

T12	R18	R19
D01	.01	.99

Explanation:

If the cause of the acute renal failure is known to be disease D01 and the test T12 is given, the probability that the test will yield test result R18 is .01 and the probability that it will yield test result R19 is .99.

T10	R04	R05	R06	R07	R08	R09	R10	R11	R12	R13	R14	R15
D01	.06	.90			.01	.01	.01		.01			
D02	.96	.01			.01	.01			.01			
D03	.98					.01			.01			
D04			.80			.03			.02		.15	
D05				.40			.20	.40				
D06	.60	.40										
D07	.05					.70			.25			
D08	.20			.10			.60	.10				
D09				.40			.20	.40				
D10	.80							.20				
D11	.05					.25			.70			
D12				.10			.20			.40		.30
D13			.40								.60	
D14				.10			.20			.30		.40






T11	R16	R17		R16	R17
D01	.05	.95	D08	.01	.99
D02	.01	.99	D09	.01	.99
D03	.99	.01	D10	.01	.99
D04	.01	.99	D11	.01	.99
D05	.01	.99	D12	.01	.99
D06	.01	.99	D13	.01	.99
D07	.15	.85	D14	.01	.99

T12	R18	R19		R18	R19
D01	.01	.99	D08	.01	.99
D02	.01	.99	D09	.01	.99
D03	.40	.60	D10	.01	.99
D04	.01	.99	D11	.01	.99
D05	.01	.99	D12	.01	.99
D06	.01	.99	D13	.01	.99
D07	.01	.99	D14	.01	.99

T13	R16	R17		R16	R17
D01	.01	.99	D08	.01	.99
D02	.01	.99	D09	.01	.99
D03	.80	.20	D10	.01	.99
D04	.01	.99	D11	.01	.99
D05	.01	.99	D12	.01	.99
D06	.01	.99	D13	.01	.99
D07	.01	.99	D14	.01	.99

T14	R20	R21		R20	R21
D01	.00	1.0	D08	.03	.97
D02	.00	1.0	D09	.85	.15
D03	.00	1.0	D10	.01	.99
D04	.00	1.0	D11	.01	.99
D05	.01	.99	D12	.01	.99
D06	.00	1.0	D13	.01	.99
D07	.00	1.0	D14	.01	.99

Date Due

    MAY 08 78 MAY 2 78 SEP 11 '78 MAY 25 80 MAY 20 81 		
---	--	--

Lib-26-67

MIT LIBRARIES



392-69

9080 003 905 509

MIT LIBRARIES



393-69

3 9080 003 905 616

MIT LIBRARIES



3 9080 003 905 608

MIT LIBRARIES



395-69

3 9080 003 874 614

MIT LIBRARIES



396-69

3 9080 003 905 574

